arXiv:1407.1268v1 [math.OC] 4 Jul 2014

ON THE CONNECTION BETWEEN THE CONJUGATE GRADIENT METHOD AND QUASI-NEWTON METHODS ON QUADRATIC PROBLEMS

Anders FORSGREN*
Tove ODLAND*

The final publication is available at Springer via
http://dx.doi.org/10.1007/s10589-014-9677-5

Abstract

It is well known that the conjugate gradient method and a quasi-Newton method, using any well-defined update matrix from the one-parameter Broyden family of updates, produce identical iterates on a quadratic problem with positive-definite Hessian. This equivalence does not hold for any quasi-Newton method. We define precisely the conditions on the update matrix in the quasi-Newton method that give rise to this behavior. We show that the crucial facts are, that the range of each update matrix lies in the last two dimensions of the Krylov subspaces defined by the conjugate gradient method and that the quasi-Newton condition is satisfied. In the framework based on a sufficient condition to obtain mutually conjugate search directions, we show that the one-parameter Broyden family is complete.

A one-to-one correspondence between the Broyden parameter and the nonzero scaling of the search direction obtained from the corresponding quasi-Newton method compared to the one obtained in the conjugate gradient method is derived. In addition, we show that the update matrices from the one-parameter Broyden family are almost always well-defined on a quadratic problem with positive-definite Hessian. The only exception is when the symmetric rank-one update is used and the unit steplength is taken in the same iteration. In this case it is the Broyden parameter that becomes undefined.

Keywords: conjugate gradient method, quasi-Newton method, unconstrained program

*Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden (andersf@kth.se, odland@kth.se). Research supported by the Swedish Research Council (VR).

1 Introduction

In this paper we examine some well-known methods used for solving unconstrained optimization problems and specifically their behavior on quadratic problems. A motivation why these problems are of interest is that the task of solving a linear system of equations $Ax = b$, with the assumption $A = A^T \succ 0$, may equivalently be considered as the one of solving an unconstrained quadratic programming problem,

$$\min_{x\in\mathbb{R}^n} q(x) = \min_{x\in\mathbb{R}^n} \frac{1}{2}x^T Hx + c^T x, \tag{QP}$$

where one lets $H = A$ and $c = -b$ to obtain the usual notation.

Given an initial guess x_0, the general idea of most methods for solving (QP) is to, in each iteration k, generate a search direction p_k and then take a steplength α_k along that direction to approach the optimal solution. For $k \geq 0$, the next iterate is hence obtained as

$$x_{k+1} = x_k + \alpha_k p_k. \tag{1}$$

The main difference between methods is the way the search direction p_k is generated. For high-dimensional problems it is preferred that only function and gradient values are used in calculations. The gradient of the objective function $q(x)$ is given by $g(x) = Hx + c$ and its value at x_k is denoted by g_k.

The research presented in this paper stems from the desire to better understand the well-known connection between the *conjugate gradient method*, henceforth CG, and *quasi-Newton methods*, henceforth QN. We are interested in determining precise conditions in association with the generation of p_k in QN such that, using exact linesearch, CG and QN will generate the same sequence of iterates as they approach the optimal solution of (QP).

CG and QN are introduced briefly in Section 2 where we also state some background results on the connection between these two methods. In Section 3 we present our results and some concluding remarks are made in Section 4.

2 Background

On (QP), naturally, we consider the use of *exact linesearch*. Then in iteration k, the optimal steplength is given by

$$\alpha_k = -\frac{p_k^T g_k}{p_k^T H p_k}. \tag{2}$$

Since $H = H^T \succ 0$, it can be shown that the descent property, $q(x_{k+1}) < q(x_k)$, holds, as long as

$$p_k^T g_k \neq 0. \tag{3}$$

Note that the usual requirement $p_k^T g_k < 0$, i.e. that is p_k is a *descent direction* with respect to $q(x)$ at x_k with steplength $\alpha_k > 0$, is included in the more general statement of (3).

Given two parallel vectors and an initial point x_k, performing exact linesearch from the initial point with respect to a given objective function along these two vectors will yield the same iterate x_{k+1}. Hence, two methods will find the same sequence of iterates if and only if the search directions generated by the two methods are parallel. Therefore, in the remainder of this paper our focus will be on parallel search directions rather than identical iterates. We will denote parallel vectors by $p \mathbin{//} p'$.

On (QP), one is generally interested in methods for which the optimal solution is found in at most n iterations. It can be shown that a sufficient property for this behavior is that the method generates search directions which are mutually conjugate with respect to H, i.e. $p_i^T H p_j = 0, \forall i \neq j$, see, e.g., [17, Chapter 5].

2.1 Conjugate gradient method

A generic way to generate conjugate vectors is by means of the conjugate Gram-Schmidt process. Given a set of linearly independent vectors $\{a_0, \ldots, a_{n-1}\}$, a set of vectors $\{p_0, \ldots, p_{n-1}\}$ mutually conjugate with respect to H can be constructed by letting $p_0 = a_0$ and for $k > 0$,

$$p_k = a_k + \sum_{j=0}^{k-1} \beta_{kj} p_j. \tag{4}$$

The values of $\{\beta_{kj}\}_{j=0}^{k-1}$ are uniquely determined in order to make p_k conjugate to $\{p_0, \ldots, p_{k-1}\}$ with respect to H. *Conjugate direction methods* is the common name for all methods which are based on generating search directions in this manner.

With the choice $a_k = -g_k$ in (4) one obtains the conjugate gradient method, CG, of Hestenes and Stiefel [11]. In effect, in CG let $p_0 = -g_0$, and for $k > 0$ the only β-value in (4) that will be non-zero is

$$\beta_{k,k-1} = \frac{p_{k-1}^T H g_k}{p_{k-1}^T H p_{k-1}}, \tag{5}$$

where one may drop the first sub-index. It can be shown that $\beta_{k-1} = g_k^T g_k / g_{k-1}^T g_{k-1}$, so that (4) may be written as

$$p_k = -g_k + \beta_{k-1} p_{k-1} = -\sum_{i=0}^{k} \frac{g_k^T g_k}{g_i^T g_i} g_i. \tag{6}$$

From the use of exact linesearch it holds that $g_k^T p_j = 0$, for all $j \leq k-1$, which implies that $g_k^T g_j = 0$, for all $j \leq k-1$, i.e. the produced gradients are orthogonal and therefore linearly independent, as required. As for any conjugate directions method using exact linesearch, the search directions of CG are descent directions. See, e.g., [20] for an intuitive introduction to the conjugate gradient method with derivation of the above relations.

In CG, one only needs to store the most recent previous search direction, p_{k-1}. This is a reduction in the amount of storage required compared to a general conjugate direction method where potentially all previous search directions are needed to compute p_k.

Although equations (1), (2), (5) and (6) give a complete description of an iteration of CG, the power and richness of the method is somewhat clouded in notation. An intuitive way to picture what happens in an iteration of CG is to describe it as a Krylov subspace method.

Definition 2.1. *Given a matrix A and a vector b the* Krylov subspace *generated by A and b is given by $\mathcal{K}_k(b, A) = span\{b, Ab, \ldots, A^{k-1}b\}$.*

Krylov subspaces are linear subspaces, which are expanding, i.e. $\mathcal{K}_1(b, A) \subseteq \mathcal{K}_2(b, A) \subseteq \mathcal{K}_3(b, A) \subseteq \ldots$, and $\dim(\mathcal{K}_k(b, A)) = k$. Given $x \in \mathcal{K}_k(b, A)$, then $Ax \in \mathcal{K}_{k+1}(b, A)$, see, e.g., [9] for an introduction to Krylov space methods. CG is a Krylov subspace method and iteration k may be formulated as the following constrained optimization problem

$$\min\ q(x),\ \text{s.t. } x \in x_0 + \mathcal{K}_{k+1}(p_0, H), \qquad (\text{CG}_k)$$

see, e.g., [17, Chapter 5]. The optimal solution of (CG_k) is x_{k+1} and the corresponding multiplier is given by $g_{k+1} = \nabla q(x_{k+1}) = Hx_{k+1} + c$. In each iteration, the dimension of the affine subspace where the optimal solution is sought increases by one. After at most n iterations the optimal solution in $\mathbb{R}^n$ is found, which will then be the optimal solution of (QP).[1]

The search direction p_k belongs to $\mathcal{K}_{k+1}(p_0, H)$, and as it is conjugate to all previous search directions with respect to H it holds that $span\{p_0, p_1, \ldots, p_k\} = \mathcal{K}_{k+1}(p_0, H)$, i.e., the search directions $p_0, \ldots, p_k$ form an H-orthogonal basis for $\mathcal{K}_{k+1}(p_0, H)$. We will henceforth refer to the search direction produced by CG, in iteration k on a given (QP), as p_k^{CG}.

Since the gradients are mutually orthogonal, and because of the relationship with the search directions in (6), it can be shown that $span\{g_0, \ldots, g_k\} = \mathcal{K}_{k+1}(p_0, H)$, i.e. the gradients form an orthogonal basis for $\mathcal{K}_{k+1}(p_0, H)$.

General conjugate direction methods can not be described as Krylov subspace methods, since in general $span\{p_0, \ldots, p_k\} \neq \mathcal{K}_{k+1}(p_0, H)$. We will use this special characteristic of CG when investigating the connection to QN.

Although our focus is quadratic programming it deserves mentioning that CG was extended to general unconstrained problems by Fletcher and Reeves [6].

2.2 Quasi-Newton methods

In QN methods the search directions are generated by solving

$$B_k p_k = -g_k, \qquad (7)$$

[1] It may happen, depending on the number of distinct eigenvalues of H and the orientation of p_0, that the optimal solution is found after less than n iterations, see, e.g., [19, Chapter 6].

in each iteration k, where the matrix B_k is chosen to be an approximation of H in some sense.[2] In this paper, we will consider *symmetric* approximations of the Hessian, i.e. $B_k = B_k^T$. It is also possible to consider unsymmetric approximation matrices, see, e.g., [12].

The first suggestion of a QN method was made by Davidon in 1959 [2], using the term *variable metric method*. In 1963, in a famous paper by Fletcher and Powell [5], Davidon's method was modified[3] and this was the starting point for making these QN methods widely known, used and studied.

We choose to work with an approximation of the Hessian rather than an approximation of the inverse Hessian, M_k, as many of the earlier papers did, e.g. [5]. Our results can however straightforwardly be derived for the inverse point of view where (7) is replaced by the equation $p_k = -M_k g_k$.

The approximation matrix B_k used in iteration k to solve for p_k is obtained by adding an *update matrix*, U_k, to the previous approximation matrix,

$$B_k = B_{k-1} + U_k. \tag{8}$$

One often considers the Cholesky factorization of B_k, then (7) can be solved in order of n^2 operations. Also, if in (8) the update matrix U_k is of low-rank, one does not need to compute the Cholesky factorization of B_k from scratch in each iteration, see, e.g., [8].

One of the most well-known update schemes is the one using update matrices from the *one-parameter Broyden family of updates* [1] described by

$$U_k = \frac{Hp_{k-1}p_{k-1}^T H}{p_{k-1}^T H p_{k-1}} - \frac{B_{k-1}p_{k-1}p_{k-1}^T B_{k-1}}{p_{k-1}^T B_{k-1} p_{k-1}} + \phi_k p_{k-1}^T B_{k-1} p_{k-1} w w^T, \tag{9}$$

with

$$w = \frac{Hp_{k-1}}{p_{k-1}^T H p_{k-1}} - \frac{B_{k-1}p_{k-1}}{p_{k-1}^T B_{k-1} p_{k-1}},$$

and where ϕ_k is a free parameter, known as the *Broyden parameter*. Equation (9) may be written more compactly, momentarily dropping all subscripts, as

$$U = \begin{pmatrix} \frac{1}{p^T Hp} Hp & \frac{1}{p^T Bp} Bp \end{pmatrix} \begin{pmatrix} p^T Hp + \varphi & -\varphi \\ -\varphi & -p^T Bp + \varphi \end{pmatrix} \begin{pmatrix} \frac{1}{p^T Hp} p^T H \\ \frac{1}{p^T Bp} p^T B \end{pmatrix}, \tag{10}$$

where $\varphi = \phi_k p^T Bp$. It is common to express the one-parameter Broyden family in terms of $y_k = g_{k+1} - g_k$ and $s_k = x_{k+1} - x_k$, see, e.g. [4], but as our interest is in the search directions we prefer the equivalent form of (9).

For all updates in this family, (8) has the property of *hereditary symmetry*, i.e. if B_{k-1} is symmetric then B_k will be symmetric. The update given by the choice

[2] The choice $B_k = H$ would give Newton's method, whereas the choice $B_k = I$ would give the steepest-descent method.

[3] "We have made both a simplification by which certain orthogonality conditions which are important to the rate of attaining the solution are preserved, and also an improvement in the criterion of convergence." [5]

$\phi_k = 0$ is known as the *Broyden-Fletcher-Goldfarb-Shanno*-update, or BFGS-update for short. For this update, when exact linesearch is used, (8) has the property of *hereditary positive definiteness*, i.e. if $B_{k-1} \succ 0$ then $B_k \succ 0$. An implication of this is that for all updates given by $\phi_k \geq 0$, when exact linesearch is used, (8) has the property of hereditary positive definiteness, see, e.g., [13, Chapter 9]. Note that there are updates in the one-parameter Broyden family for which (8) does not have this property.

2.3 Background results

In [3], Dixon has shown that on any smooth function, using perfect linesearch[4], the one-parameter Broyden family gives rise to parallel search directions. On (QP), these search directions will in addition be mutually conjugate with respect to H, see, e.g., [12].

On (QP) it is well-known that these conjugate search directions, generated by the one-parameter Broyden family, will be parallel to those of CG, i.e. $p_k \,//\, p_k^{CG}$, for all k. See, e.g., [4, 12, 15]. In particular Theorem 3.4.2 on page 65 of [4] states that CG and QN using well-defined update matrices from the one-parameter Broyden family generate identical iterates. Note that this connection between CG and QN does not hold for general convex functions.

In this paper we approach the well-known connection between QN and CG from another perspective. The main question handled in this paper is: when solving (QP), what are the precise conditions on B_k and U_k such that $p_k \,//\, p_k^{CG}$ is obtained? We provide an answer by turning our attention to the search directions defined by CG and the Krylov subspaces they span. In Proposition 3.3 we state explicit conditions on U_k and in Theorem 3.6 we extend the well-known connection between CG and QN as we show that, under a sufficient condition to generate conjugate search directions, no other update matrices than those in the one-parameter Broyden family will make $p_k \,//\, p_k^{CG}$.

3 Results

As a reminder for the reader we will, in the following proposition, state the necessary and sufficient conditions for a vector p_k to be parallel to the vector p_k^{CG} given that $p_i \,//\, p_i^{CG}$, for all $i \leq k-1$. Given x_0, one may calculate $g_0 = Hx_0 + c$.

Proposition 3.1. *Let $p_0 = p_0^{CG} = -g_0$ and $p_i \,//\, p_i^{CG}$, for all $i \leq k-1$. Assume $p_k \neq 0$, then $p_k \,//\, p_k^{CG}$ if and only if*

(i) $p_k \in \mathcal{K}_{k+1}(p_0, H)$, *and,*

(ii) $p_k^T H p_i = 0, \quad \forall i \leq k-1.$

Proof. For the sake of completeness we include the proof.

[4] A generalization of exact linesearch for general smooth functions, see [3].

Necessary: Suppose p_k is parallel to p_k^{CG}. Then $p_k = \delta_k p_k^{CG}$, for an arbitrary nonzero scalar δ_k. As p_k^{CG} satisfies (i), it holds that

$$p_k = \delta_k p_k^{CG} \in \mathcal{K}_{k+1}(p_0, H),$$

since $\mathcal{K}_{k+1}(p_0, H)$ is a linear subspace. And as p_k^{CG} satisfies (ii), it follows that,

$$p_k^T H p_i = \delta_k \delta_i (p_k^{CG})^T H p_i^{CG} = 0, \quad \forall i \leq k-1.$$

Sufficient: Suppose p_k satisfies (i) and (ii). The set of vectors $\{p_0^{CG}, \ldots, p_k^{CG}\}$ form an H-orthogonal basis for the space $\mathcal{K}_{k+1}(p_0, H)$ and the set of vectors $\{p_0, \ldots, p_{k-1}\}$ form an H-orthogonal basis for the space $\mathcal{K}_k(p_0, H)$. Since p_k satisfies (i) and (ii) it must hold that $p_k \mathbin{//} p_k^{CG}$. □

These necessary and sufficient conditions will serve as a foundation for the rest of our results. We will determine conditions, first on B_k, and second on U_k used in iteration k of QN, in order for $p_k \mathbin{//} p_k^{CG}$.

Given the current iterate x_k, one may calculate $g_k = Hx_k + c$. Hence, in iteration k of QN, p_k is determined from (7) and depends entirely on the choice of B_k. We make the assumption that (7) is compatible, i.e. that a solution exists. A well-known sufficient condition for p_k to satisfy (ii) of Proposition 3.1 is for B_k to satisfy

$$B_k p_i = H p_i, \quad \forall i \leq k-1. \tag{11}$$

This condition, known as the *hereditary condition*, will be used in the following results, which implies that the stated conditions on B_k and later U_k will be *sufficient* conditions, and not necessary and sufficient as those in Proposition 3.1. All B_k satisfying (11) can be seen as defining different conjugate direction methods, but only certain choices of B_k will generate conjugate directions parallel to those of CG.

Assume that B_k is constructed as,

$$B_k = I + V_k, \tag{12}$$

i.e. an identity matrix plus a matrix V_k.[5] We make no assumptions on V_k except symmetry.

Let $B_0 = I$, then $p_0 = p_0^{CG} = -g_0$. Given $p_i \mathbin{//} p_i^{CG}$, for all $i \leq k-1$, the following lemma gives sufficient conditions on B_k in order for $p_k \mathbin{//} p_k^{CG}$.

Lemma 3.2. *Let $B_0 = I$, so that $p_0 = p_0^{CG} = -g_0$. Assume that $p_i \mathbin{//} p_i^{CG}$, for all $i \leq k-1$, and that a solution p_k is obtained from (7). Then $p_k \mathbin{//} p_k^{CG}$ if B_k satisfies*

(i) $\mathcal{R}(V_k) \subseteq \mathcal{K}_{k+1}(p_0, H)$, and,

(ii) $B_k p_i = H p_i, \quad \forall i \leq k-1$,

where $\mathcal{R}(V_k) = \{y \ : \ y = V_k x\}$, the range-space *of the matrix V_k.*

[5] Any symmetric matrix can be expressed like this.

Proof. First we show that p_k satisfies condition (i) of Proposition 3.1. If (12) is inserted into (7) one obtains,

$$\big(I + V_k\big)p_k = -g_k,$$

so that

$$p_k = -g_k - V_k p_k.$$

Since $g_k \in \mathcal{K}_{k+1}(p_0, H)$, and $\mathcal{R}(V_k) \subseteq \mathcal{K}_{k+1}(p_0, H)$, it holds that $p_k \in \mathcal{K}_{k+1}(p_0, H)$. Hence, p_k satisfies condition (i) of Proposition 3.1. Secondly, p_k satisfies condition (ii) of Proposition 3.1 since condition (ii) of this lemma is identical to (11).

Hence, $p_k \;/\!/\; p_k^{CG}$ by Proposition 3.1. □

As B_k is updated according to (8), one would prefer to have conditions, in iteration k, on the update matrix U_k instead of on the entire matrix B_k. Therefore, we now modify Lemma 3.2 by noting that equation (12) may be stated as

$$B_k = B_{k-1} + U_k = I + V_{k-1} + U_k. \tag{13}$$

We make no assumptions on U_k except symmetry. Note that one may split (11) as

$$B_k p_i = H p_i, \quad \forall i \leq k-2, \tag{14}$$

$$B_k p_{k-1} = H p_{k-1}. \tag{15}$$

Equation (15) is known as the *quasi-Newton condition*.[6] Using (8) one can reformulate (14) and (15) in terms of U_k, see (ii) and (iii) of the following proposition.

Given B_{k-1} that satisfies condition (i) of Lemma 3.2, and given $p_i \;/\!/\; p_i^{CG}$, for all $i \leq k-1$, the following proposition gives sufficient conditions on U_k in order for $p_k \;/\!/\; p_k^{CG}$.

Proposition 3.3. *Let $B_0 = I$, so that $p_0 = p_0^{CG} = -g_0$. Assume that $p_i \;/\!/\; p_i^{CG}$, for all $i \leq k-1$, that B_{k-1} satisfies condition (i) of Lemma 3.2, that a solution p_k is obtained from* (7) *and that B_k is obtained from* (13)*. Then $p_k \;/\!/\; p_k^{CG}$ if U_k satisfies*

(i) $\mathcal{R}(U_k) \subseteq \mathcal{K}_{k+1}(p_0, H)$,

(ii) $U_k p_i = 0, \quad \forall i \leq k-2$, *and,*

(iii) $U_k p_{k-1} = (H - B_{k-1})p_{k-1}$.

Proof. The proof is identical to the one of Lemma 3.2. If (13) is inserted into (7) one obtains

$$p_k = -g_k - V_{k-1}p_k - U_k p_k.$$

By condition (i) of Lemma 3.2, $\mathcal{R}(V_{k-1}) \subseteq \mathcal{K}_k(p_0, H)$, and since $\mathcal{K}_k(p_0, H) \subseteq \mathcal{K}_{k+1}(p_0, H)$, it follows that $V_{k-1}p_k \in \mathcal{K}_{k+1}(p_0, H)$. Therefore, since

[6] In the literature, much emphasis is placed on that updates should satisfy (15). This condition alone is not a sufficient condition on B_k to give conjugate directions.

$g_k \in \mathcal{K}_{k+1}(p_0, H)$ and $\mathcal{R}(U_k) \subseteq \mathcal{K}_{k+1}(p_0, H)$, it holds that $p_k \in \mathcal{K}_{k+1}(p_0, H)$. Hence, p_k satisfies condition (i) of Proposition 3.1.

Conditions (ii) and (iii) of this lemma are merely a reformulation of (11), hence p_k satisfies condition (ii) of Proposition 3.1.

Hence, $p_k \mathbin{//} p_k^{CG}$ by Proposition 3.1. □

The assumption of the previous proposition, that B_{k-1} is chosen to satisfy condition (i) of Lemma 3.2, will be satisfied if, in each iteration k, the update matrix U_k is chosen according to conditions (i)-(iii) of Proposition 3.3. This is summarized in the following corollary.

Corollary 3.4. *If in each iteration k the update matrix U_k is chosen to satisfy conditions (i)-(iii) of Proposition 3.3, then $p_k \mathbin{//} p_k^{CG}$ for all k.*

Next we state a result, which will be needed in our further investigation, and holds for any update scheme of B_k, which generates search directions parallel to those of CG.

Proposition 3.5. *If $p_k \mathbin{//} p_k^{CG}$ for all k and a solution p_k obtained from* (7)*, then $p_k^T B_k p_k \neq 0$, unless $B_k p_k = -g_k = 0$.*

Proof. Since $p_k = \delta_k p_k^{CG}$, for some non-zero scalar δ_k, and since p_k^{CG} is a descent direction it holds that

$$p_k^T B_k p_k = -p_k^T g_k = -\delta_k (p_k^{CG})^T g_k \neq 0,$$

for some non-zero scalar δ_k, unless $g_k = 0$. □

Note that this implies that (3) is satisfied for any QN method using an update scheme that generates search directions that are parallel to those of CG. Also, assuming $g_k \neq 0$, it implies that the fraction $\frac{1}{p_k^T B_k p_k}$ is well-defined.

3.1 Update matrices defined by Proposition 3.3

Having stated precise conditions on U_k in Proposition 3.3 we now turn to look at what these conditions imply in terms of actual update matrices. In Theorem 3.6, we show that the conditions on U_k in Proposition 3.3 are equivalent to the matrix U_k belonging to the one-parameter Broyden family, (10). This implies that, under the sufficient condition (11), there are no other update matrices, even of higher rank, that make $p_k \mathbin{//} p_k^{CG}$.

The if-direction of Theorem 3.6, the fact that the one-parameter Broyden family satisfies the conditions of Proposition 3.3 is straightforward. However, the only if-directions shows that there are no update matrices outside this family that satisfy the conditions of Proposition 3.3.

Theorem 3.6. *Assume $g_k \neq 0$. A matrix U_k satisfies (i)-(iii) of Proposition 3.3 if and only if U_k can be expressed according to* (10)*, the one-parameter Broyden family, for some ϕ_k.*

Proof. Note that since $\mathcal{K}_{k-1}(p_0, H) = span\{p_0, \ldots, p_{k-2}\}$, condition (ii) of Proposition 3.3 can be stated as

$$\mathcal{K}_{k-1}(p_0, H) \subseteq \mathcal{N}(U_k),$$

where $\mathcal{N}(U_k) = \{x \,:\, U_k x = 0\}$, the *null-space* of U_k. This implies that $\dim(\mathcal{N}(U_k)) \geq k-1$. Since U_k is symmetric and applying condition (i) of Proposition 3.3 it follows that

$$\mathcal{R}(U_k) \subseteq \mathcal{K}_{k-1}(p_0, H)^{\perp} \cap \mathcal{K}_{k+1}(p_0, H) = span\{g_{k-1}, g_k\},$$

and that $\dim(\mathcal{R}(U_k^T)) = \dim(\mathcal{R}(U_k)) \leq 2$.

Hence, one may write a general U_k that satisfies the conditions (i) and (ii) of Proposition 3.3 as

$$U_k = \begin{pmatrix} g_{k-1} & g_k \end{pmatrix} \begin{pmatrix} m_{1,1} & m_{1,2} \\ m_{1,2} & m_{2,2} \end{pmatrix} \begin{pmatrix} g_{k-1}^T \\ g_k^T \end{pmatrix}. \tag{16}$$

Note that due to the linear relationship between $\{B_{k-1}p_{k-1}, Hp_{k-1}\}$ and $\{g_{k-1}, g_k\}$ it holds that

$$g_{k-1} = -B_{k-1}p_{k-1}, \quad g_k = \alpha_{k-1}Hp_{k-1} + g_{k-1}.$$

Hence, since $\alpha_{k-1} = (p_{k-1}^T B_{k-1} p_{k-1})/(p_{k-1}^T H p_{k-1})$ by the exact linesearch, dropping all $'k-1'$-subscripts and letting g_k be represented by g_+, we obtain

$$\begin{pmatrix} g^T \\ g_+^T \end{pmatrix} = p^T B p \begin{pmatrix} 0 & -1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} \frac{1}{p^T H p} p^T H \\ \frac{1}{p^T B p} p^T B \end{pmatrix}. \tag{17}$$

We may therefore rewrite (16) as

$$U_k = \begin{pmatrix} \frac{1}{p^T H p} Hp & \frac{1}{p^T B p} Bp \end{pmatrix} \begin{pmatrix} \hat{m}_{1,1} & \hat{m}_{1,2} \\ \hat{m}_{1,2} & \hat{m}_{2,2} \end{pmatrix} \begin{pmatrix} \frac{1}{p^T H p} p^T H \\ \frac{1}{p^T B p} p^T B \end{pmatrix}. \tag{18}$$

Now imposing condition (iii) of Proposition 3.3, (18) yields

$$U_k p = \begin{pmatrix} \frac{1}{p^T H p} Hp & \frac{1}{p^T B p} Bp \end{pmatrix} \begin{pmatrix} \hat{m}_{1,1} & \hat{m}_{1,2} \\ \hat{m}_{1,2} & \hat{m}_{2,2} \end{pmatrix} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = Hp - Bp. \tag{19}$$

A combination of Proposition 3.5 and (17) shows that Hp and Bp are linearly independent. Hence, (19) implies that

$$\hat{m}_{1,1} + \hat{m}_{1,2} = p^T H p, \quad \hat{m}_{1,2} + \hat{m}_{2,2} = -p^T B p.$$

With $\hat{m}_{1,2} = -\varphi$, then

$$\hat{m}_{1,1} = p^T H p + \varphi, \quad \hat{m}_{2,2} = -p^T B p + \varphi.$$

Substituting into (18), we obtain (10) with the scaling $\varphi = \phi_k p^T B p$. This completes the proof. □

The precise conditions of Proposition 3.3 are equivalent to choosing U_k from the one-parameter Broyden family. As seen in the proof, the conditions (i)-(iii) of Proposition 3.3 may be expressed as

(a) $\mathcal{R}(U_k) \subseteq span\{g_{k-1}, g_k\} = \mathcal{K}_{k-1}(p_0, H)^\perp \cap \mathcal{K}_{k+1}(p_0, H)$, and,

(b) $U_k p_{k-1} = (H - B_{k-1})p_{k-1}$.

Note that after we impose (a) we have three degrees of freedom, $\hat{m}_{1,1}$, $\hat{m}_{1,2}$ and $\hat{m}_{2,2}$. Then after imposing (b), the *quasi-Newton condition*, we are left with one degree of freedom ϕ_k.

What is often mentioned as a feature of the one-parameter Broyden family, using information only from the current and previous iteration when forming U_k, is in fact a condition that guarantees the equivalence to CG. Also, the fact that the update matrices in (10) are of rank at most two is a consequence of satisfying this condition. This is what distinguishes the one-parameter Broyden family of updates from any B_k satisfying (11).

We stress that the fact that CG and QN, using a well-defined update matrix from the one-parameter Broyden family, generates parallel search directions and hence identical iterates is well-known. By Proposition 3.3 and Theorem 3.6 we may draw the additional conclusion that, under the sufficient condition (11), there are no other update matrices, even of higher rank, that make $p_k \,/\!/\, p_k^{CG}$.

3.2 Relation between δ_k and ϕ_k

Next we derive a relation between the free parameter ϕ_k, the Broyden parameter used in U_k when forming B_k according to (8), and the non-zero parameter δ_k in $p_k = \delta_k p_k^{CG}$. In the following lemma we state an alternative way to express U_k belonging to the one-parameter Broyden family (10).

Lemma 3.7. *An update matrix U_k from the one-parameter Broyden family* (10) *can be expressed on the form*

$$U_k = \begin{pmatrix} g_{k-1} & g_k \end{pmatrix} \begin{pmatrix} \frac{p^THp}{(p^TBp)^2} - \frac{1}{p^TBp} & -\frac{p^THp}{(p^TBp)^2} \\ -\frac{p^THp}{(p^TBp)^2} & \frac{p^THp}{(p^TBp)^2} + \frac{\phi_k}{p^TBp} \end{pmatrix} \begin{pmatrix} g_{k-1}^T \\ g_k^T \end{pmatrix}, \tag{20}$$

where all omitted subscripts are $'k-1'$.

Proof. Reversing relation (17) gives

$$\begin{pmatrix} \frac{1}{p^THp}p^TH \\ \frac{1}{p^TBp}p^TB \end{pmatrix} = \frac{1}{p^TBp} \begin{pmatrix} -1 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} g_{k-1}^T \\ g_k^T \end{pmatrix}. \tag{21}$$

If (21) is inserted in (10) with $\varphi = \phi_k p^TBp$, then (20) is obtained. □

In the following proposition we show the one-to-one correspondence between the two parameters ϕ_k and δ_k,

Proposition 3.8. *Let $B_0 = I$, so that $p_0 = p_0^{CG} = -g_0$. Assume that $p_i \,/\!/\, p_i^{CG}$, for all $i \leq k-1$, that B_{k-1} satisfies condition (i) of Lemma 3.2, that a solution p_k is obtained from* (7)*, that B_k is obtained from* (13) *and that U_k belongs to the one-parameter Broyden family,* (10)*, for some ϕ_k, where $\phi_k \neq -(p_{k-1}^TB_{k-1}p_{k-1})/(g_k^Tg_k)$.*

Then the non-zero parameter δ_k in $p_k = \delta_k p_k^{CG}$ is given by

$$\delta_k(\phi_k) = \frac{1}{1 + \phi_k \frac{g_k^T g_k}{p_{k-1}^T B_{k-1} p_{k-1}}}. \tag{22}$$

Proof. Omitting all $'k-1'$-subscripts throughout the proof. Our assumption is that p_k is given by

$$p_k = -g_k - V_{k-1}p_k - U_k p_k.$$

By Theorem 3.6 $p_k \;//\; p_k^{CG}$ hence,

$$\delta_k p_k^{CG} = -g_k - V_{k-1}\delta_k p_k^{CG} - U_k \delta_k p_k^{CG}.$$

If the above expression is projected onto g_k it holds that

$$g_k^T(\delta_k p_k^{CG}) = -g_k^T g_k - \delta_k g_k^T V_{k-1} p_k^{CG} - \delta_k g_k^T U_k p_k^{CG} = -g_k^T g_k - 0 - \delta_k g_k^T U_k p_k^{CG},$$

since by condition (i) of Lemma 3.2, $\mathcal{R}(V_{k-1}) \subseteq \mathcal{K}_k(p_0, H) = span\{g_0, g_1, \ldots, g_{k-1}\}$. Using (20) for U_k it follows that

$$\delta_k g_k^T p_k^{CG} = -g_k^T g_k - \delta_k g_k^T g_k \begin{pmatrix} -\frac{p^T H p}{(p^T B p)^2} & \frac{p^T H p}{(p^T B p)^2} + \frac{\phi_k}{p^T B p} \end{pmatrix} \begin{pmatrix} g_{k-1}^T p_k^{CG} \\ g_k^T p_k^{CG} \end{pmatrix}.$$

From (6) it follows that $g_i^T p_k^{CG} = -g_k^T g_k$, for all $i \leq k$, which implies that

$$-\delta_k g_k^T g_k = -g_k^T g_k - \delta_k g_k^T g_k \begin{pmatrix} -\frac{p^T H p}{(p^T B p)^2} & \frac{p^T H p}{(p^T B p)^2} + \frac{\phi_k}{p^T B p} \end{pmatrix} \begin{pmatrix} -g_k^T g_k \\ -g_k^T g_k \end{pmatrix},$$

dividing on both sides with $-g_k^T g_k$ yields

$$\delta_k = 1 - \delta_k g_k^T g_k \big(-\frac{p^T H p}{(p^T B p)^2} + \frac{p^T H p}{(p^T B p)^2} + \frac{\phi_k}{p^T B p} \big) = 1 - \delta_k g_k^T g_k \frac{\phi_k}{p^T B p},$$

which implies (22).

□

Note that this result implies that for $\phi_k = 0$ we get $\delta_k = 1$, i.e. $p_k^{BFGS} = p_k^{CG}$ as long as ϕ_i, for all $i \leq k-1$, are well-defined[7]. In [15], Nazareth derives the relation between p_k^{BFGS} and p_k^{CG} using induction.

Note that substituting the limit $\phi_k \to -(p_{k-1}^T B_{k-1} p_{k-1})/(g_k^T g_k)$ into (22) for $g_k \neq 0$ yields $\delta_k \to \infty$. This limit of ϕ_k is called a *degenerate value* since it makes B_k singular, which is a contradiction to our assumption that $B_k p_k = -g_k$ is compatible. For all other values of ϕ_k, we get a one-to-one correspondence between ϕ_k and δ_k and this implies that our assumption that $B_k p_k = -g_k$ is compatible actually implies that the solution p_k is unique.

The assumption of Proposition 3.8, that B_{k-1} is chosen to satisfy condition (i) of Lemma 3.2, will be satisfied if, in each iteration k, the update matrix U_k belongs to the one-parameter Broyden family, (10), for some ϕ_k.

[7] See the next paragraph and Section 3.3.

3.3 Remarks on when the one-parameter Broyden family is well-defined

There are several ways in which an update scheme using the one-parameter Broyden family may become not well-defined. We have already mentioned the degenerate value that makes B_k singular in the previous section.

In addition to the above, the matrix U_k may itself become not well-defined. For all well-defined values of ϕ_k it holds U_k given by (10) is not well-defined if and only if $p_{k-1}^T B_{k-1} p_{k-1} = 0$ and $B_{k-1} p_{k-1} \neq 0$. It is clear that requiring B_{k-1} to be definite (positive or negative) is sufficient to avoid U_k being not well-defined. However, from Proposition 3.5 it follows that, on (QP), $p_{k-1}^T B_{k-1} p_{k-1} = 0$ and $B_{k-1} p_{k-1} \neq 0$ does not occur for any update scheme which generates search directions parallel to those of CG.[8]

Hence, on (QP), in order for U_k, given by (10), to become not well-defined, the undefinedness must enter in the Broyden parameter ϕ_k. It is well-known that the symmetric rank-one update, SR1, may become not well-defined on (QP), see, e.g. [13, Chapter 9]. SR1 is uniquely determined by (15)[9] and the Broyden parameter for SR1 is given by $\phi_k^{SR1} = (p_{k-1}^T H p_{k-1})/(p_{k-1}^T (H - B_{k-1}) p_{k-1})$, an expression that becomes undefined for $p_{k-1}^T (H - B_{k-1}) p_{k-1} = 0$. Note that this is equivalent to $\alpha_{k-1} = (p_{k-1}^T B_{k-1} p_{k-1})/(p_{k-1}^T H p_{k-1}) = 1$.

Hence, we may summarize our remarks on when U_k given by the one-parameter Broyden family, (10), is well-defined in the following corollary.

Corollary 3.9. *Unless $\phi_k = \phi_k^{SR1}$ in* (10) *with the unit steplength taken in the same iteration, then U_k defined by* (10)*, the one-parameter Broyden family, is always well-defined on* (QP)*.*

Hence, taking the unit steplength is an indication that one needs to choose a different update scheme than SR1 when forming B_k. We therefore stress that hereditary (positive or negative) definiteness is not a necessary property for the update matrices U_k to be well-defined when solving (QP).

4 Conclusions

The main result of this paper are the precise conditions on the update matrix U_k stated in Proposition 3.3. In addition, from Theorem 3.6, we draw the conclusion that, in the framework where we use the sufficient condition (11) to guarantee conjugacy of the search directions, the update schemes in QN that give parallel search directions to those of CG are completely described by the one-parameter Broyden family. Hence, we are able to extend the well-known connection between CG and QN, e.g., given in [4]. We show that, under the sufficient condition (11), there are no other update matrices, even of higher rank, that make $p_k \mathbin{/\!/} p_k^{CG}$.

[8] For general functions we may have $p_{k-1}^T B_{k-1} p_{k-1} = 0$ and $B_{k-1} p_{k-1} \neq 0$, and in [7] the values of ϕ_k that give rise to this situation are characterized and also termed degenerate values.

[9] An illustration that the conditions (i)-(iii) of Proposition 3.3 are indeed only sufficient.

It seems as it may be the sufficient requirement to get conjugate directions, (11) that limits the freedom when choosing U_k to the one-parameter Broyden family of updates as shown by the above results. Since this condition is only sufficient, one may still pose the question if there are other update schemes for QN that yield the same sequence of iterates as CG on (QP). We believe that in order to understand this possible limitation it will be necessary to obtain a deeper understanding of CG.

We derive a one-to-one correspondence between the Broyden parameter ϕ_k and the scaling δ_k. This relation implies that the assumption that (7) is compatible implies uniqueness of the solution p_k. If the degenerate value is used B_k becomes singular.

We are also able to make the remark that the update matrices belonging to the one-parameter Broyden family is always well-defined on (QP), unless the steplength is of unit length and in the same iteration the rank-one update is used. In this case it is the Broyden parameter ϕ_k that becomes undefined.

In this paper we have focused on quadratic programming. Besides being important in its own right, it is also a highly important as a subproblem when solving unconstrained optimization problems. For a survey on methods for unconstrained optimization see, e.g., [18]. Also, it deserves mentioning that work has been done on QN update schemes for general unconstrained optimization considering (11) for only $j = k - 1$ and $j = k - 2$, deriving an update scheme that satisfies the quasi-Newton condition and has a minimum violation of it for the previous step, see [14, 16].

A further motivation for this research in this paper is that the deeper understanding of what is important in the choice of U_k could be implemented in a limited-memory QN method. I.e., can one choose which columns to save based on some other criteria than just picking the most recent ones? See, e.g., [17, Chapter 9], for an introduction to limited-memory QN methods.

Finally, it should be pointed out that the discussion of this paper is limited to exact arithmetic. Even in cases where CG and QN generate identical iterates in exact arithmetic, the difference between numerically computed iterates by the two methods may be quite large, see, e.g. [10].

Acknowledgements:

We thank the referees and the associate editor for their insightful comments and suggestions which significantly improved the presentation of this paper.

References

[1] Broyden, C. G. Quasi-Newton methods and their application to function minimisation. *Math. Comp. 21* (1967), 368–381.

[2] Davidon, W. C. Variable metric method for minimization. *SIAM J. Optim. 1*, 1 (1991), 1–17.

[3] Dixon, L. C. W. Quasi-Newton algorithms generate identical points. *Math. Programming 2* (1972), 383–387.

[4] Fletcher, R. *Practical methods of optimization*, second ed. A Wiley-Interscience Publication. John Wiley & Sons Ltd., Chichester, 1987.

[5] Fletcher, R., and Powell, M. J. D. A rapidly convergent descent method for minimization. *Comput. J. 6* (1963/1964), 163–168.

[6] Fletcher, R., and Reeves, C. M. Function minimization by conjugate gradients. *Comput. J. 7* (1964), 149–154.

[7] Fletcher, R., and Sinclair, J. W. Degenerate values for Broyden methods. *J. Optim. Theory Appl. 33*, 3 (1981), 311–324.

[8] Gill, P. E., Murray, W., and Wright, M. H. *Practical optimization*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1981.

[9] Gutknecht, M. H. A brief introduction to Krylov space methods for solving linear systems. In *Frontiers of Computational Science* (2007), Springer-Verlag Berlin, pp. 53–62. International Symposium on Frontiers of Computational Science, Nagoya Univ, Nagoya, Japan, Dec 12-13, 2005.

[10] Hager, W. W., and Zhang, H. The limited memory conjugate gradient method. *SIAM J. Optim. 23*, 4 (2013), 2150–2168.

[11] Hestenes, M. R., and Stiefel, E. Methods of conjugate gradients for solving linear systems. *J. Research Nat. Bur. Standards 49* (1952), 409–436 (1953).

[12] Huang, H. Y. Unified approach to quadratically convergent algorithms for function minimization. *J. Optimization Theory Appl. 5* (1970), 405–423.

[13] Luenberger, D. G. *Linear and nonlinear programming*, second ed. Addison-Wesley Pub Co, Boston, MA, 1984.

[14] Mifflin, R. B., and Nazareth, J. L. The least prior deviation quasi-Newton update. *Math. Programming 65*, 3, Ser. A (1994), 247–261.

[15] Nazareth, L. A relationship between the BFGS and conjugate gradient algorithms and its implications for new algorithms. *SIAM J. Numer. Anal. 16*, 5 (1979), 794–800.

[16] Nazareth, L. An alternative variational principle for variable metric updating. *Math. Programming 30*, 1 (1984), 99–104.

[17] Nocedal, J., and Wright, S. J. *Numerical optimization*. Springer Series in Operations Research. Springer-Verlag, New York, 1999.

[18] Powell, M. J. D. Recent advances in unconstrained optimization. *Math. Programming 1*, 1 (1971), 26–57.

[19] Saad, Y. *Iterative methods for sparse linear systems*, second ed. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2003.

[20] Shewchuk, J. R. An introduction to the conjugate gradient method without the agonizing pain. Tech. rep., Pittsburgh, PA, USA, 1994.